arXiv:1703.06843v1 [cs.SI] 20 Mar 2017

The Role of Network Analysis in Industrial and Applied Mathematics

Mason A. Porter[1,2,3] and Sam D. Howison[2]
[1]Department of Mathematics, University of California, Los Angeles, 90095, USA
[2]Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK
[3]CABDyN Complexity Centre, University of Oxford, Oxford OX1 1HP, UK

March 21, 2017

Abstract

Many problems in industry — and in the social, natural, information, and medical sciences — involve discrete data and benefit from approaches from subjects such as network science, information theory, optimization, probability, and statistics. Because the study of networks is concerned explicitly with connectivity between different entities, it has become very prominent in industrial settings, and this importance has been accentuated further amidst the modern data deluge. In this article, we discuss the role of network analysis in industrial and applied mathematics, and we give several examples of network science in industry.

1 Introduction

Mathematics has played a vital role in industry for a long time. From the mixing of fluids to produce an ideal bar of chocolate to the study of gasoline emissions in vehicles, wonderful problems in continuous mathematics — typically framed in terms of ordinary and partial differential equations — have arisen from industrial problems; see the online repository [1]. They have increasingly complemented the equally wonderful problems posed by applications in science and engineering. The applied mathematics curricula (and subjects studied by the academic staff) of universities such as University of Oxford and University of Cambridge reflect this bias toward continuum mechanics.

One of the most prominent approaches to applied mathematics, which has been particularly prominent in the United Kingdom, is "physical applied mathematics". This entails starting from a problem (most traditionally from the physical sciences); using basic physical (or biological or chemical) principles and, more generally, using relevant domain knowledge to derive equations of motion (most often in the form of ordinary or partial differential equations) and boundary and/or initial conditions; simplifying the equations to make them

mathematically tractable (e.g., perhaps some terms can be ignored); studying the equations both computationally and (often approximately, such as with asymptotics and perturbation theory) with a wide variety of mathematical tools; comparing the numerical solutions (ideally of both the simplified equations of motion and, if possible, the "original" equations) with approximate analytical solutions or qualitative behavior revealed through procedures like a dynamical-systems analysis; comparing these results with controlled experiments; and, where possible, comparing the experimental results with natural or industrial phenomena in more realistic settings. Textbooks such as [40, 20] describe these ideals. Through making comparisons, one also refines one's assumptions (e.g., maybe some of them turn out to be inappropriate), adjusts models, refines experiments, and so on. A final crucial step is to interpret the results of the mathematical and numerical studies in a way that engages seriously with the original problem. A problem's stakeholders must learn something from the mathematical efforts, and such stakeholders (whether they are scientists in other departments, people who work in industry or government, or someone else) often collaborate directly on the problem, and at minimum they need to be consulted early and often, as they offer domain knowledge.

Good physical applied mathematics can start from potentially any type of problem, including fluid, solid, or granular phenomena in nature; observations of biological systems; observations of human or animal behavior; physical, behavioral, or other phenomena in industry; and much more. Industrial problems (and others as well) typically start out in woolly form, and the key challenge for applied mathematicians is constructing a concrete, tractable mathematical problem whose solution (perhaps in approximate or numerical form) can yield important insights about the original problem or phenomenon. This is the art of "mathematical modeling," and in industrial mathematics, it entails taking a physical-applied-mathematics approach to problems that arise from industry. It is APPLIED mathematics rather than simply applied MATHEMATICS, as it is crucial to engage very seriously with applications.

This approach, for which UK and European study groups with industry have made pioneering contributions [1], also applies to problems and data that take discrete forms. It has long been true that many problems in industry include discrete data and benefit from approaches that incorporate topics from subjects such as network science, information theory, optimization, probability, and statistics. For example, solving problems in optimization is crucial for assembly lines, and the famous traveling salesperson problem (TSP) has an undeniably practical origin [13]. Amidst the modern data deluge, the importance of discrete data and associated approaches has reached new heights. Social media, which are now incorporated throughout industry, involve interactions between entities; radio-frequency identification (RFID) device data, with finite resolution, track the movements of people in cities and stores; people have associated metadata that describe their characteristics using discrete (categorical or ordinal) variables, and so on. Because network science is concerned explicitly with connectivity between (and among) different entities, it has become very prominent in industrial settings, and this importance has been accentuated further

by the newfound wealth of data.

Many discrete data sets, and problems in which some kind of interactions or coupling are relevant, benefit from the study of networks, which has become one of the most prominent areas of applied mathematics, physics, computer science, and other disciplines [48, 52]. The basic idea is to use the established approach of physical applied mathematics, but now the problems need not be physical (or from other traditional domains), and in particular are often discrete in nature and/or involve copious amounts of data. It is desirable to combine network science, "the study of connectivity", with an applied mathematics philosophy, which has been enormously successful in collaborations with industry [2].[1] For a recent collection of modeling efforts involving networks (including in industry), see the December 2016 special issue of *European Journal of Applied Mathematics* [52], which one of us (MAP) edited along with Ginestra Bianconi of Queen Mary University of London. For some more specific examples, see the special issue of *Royal Society Open Science* on urban analytics [27].

In this article, we discuss the role of network analysis in industrial and applied mathematics, and we give several examples of network science in industry. In Section 2, we discuss network modeling and relate it to traditional topics in physical applied mathematics. In Section 3, we discuss applications of mathematics to the social sciences. In Section 4, we discuss network science in industrial settings. We conclude in Section 5.

2 Network Modeling

To appreciate how to use a physical-applied-mathematical approach in the study of networks, one first needs to have some idea of what a network is. In its broadest form, a network is a representation of the connectivity patterns and strengths in a complex systems of interacting entities [48, 63, 35, 33, 60]. Most traditionally, a network is represented mathematically as a graph G, which consists of a set V of "nodes" (or "vertices") that encode entities and a set $E \subseteq V \times V$ of "edges" (or "links" or "ties") that encode the interactions between those entities. However, the term "network" is more general than a graph, as a network can encompass connections among an arbitrary number of entities, can have nodes and/or edges that change in time, can include multiple types of edges, can often have associated dynamical processes both on the networks and of the networks, and so on. Associated with a graph is an "adjacency matrix" $\mathbf{A}$, where (if one does not include a value to model the strength of a connection) an entry $a_{ij} = 1$ indicates the presence of an edge connecting entity i to entity j directly, while $a_{ij} = 0$ indicates its absence. That is, when $a_{ij} = 1$, nodes i and j are "adjacent" to each other, and the associated edge is "incident" to each of the two nodes. To consider edges with different levels of importance, one can assign a weight — typically a nonnegative real number, although there are many situations, such as in the study of international relations or social

[1] The paper [26] in this special issue gives an example application of network analysis in industry.

interactions, in which negative values can be appropriate — so that the entry w_{ij} of a "weighted adjacency matrix" (or "weight matrix") $\mathbf{W}$ represents the weight of the connection between nodes i and j. In most applications, a high value of w_{ij} represents a strong connection between entities i and j, though sometimes (e.g., for applications in transportation) one might instead have a distance matrix, and then elements with smaller values represent stronger connections. Such data can arise, for example, in the form of physical distances (e.g., road networks) or from measurements of empirical or expected travel time between a pair of locations (e.g., using Oyster card data from Transport for London). Connections in a network can be directed; in that case, a_{ij} need not equal a_{ji} (and w_{ij} need not equal w_{ji}), leading (aside from fully reciprocal situations, which are rare) to asymmetric adjacency and weight matrices.

The study of networks incorporates tools from a wide variety of areas — such as graph theory (of course), computational linear algebra, dynamical systems, optimization, statistical physics, probability, statistics, and more — and is important for applications in just about any area that one might imagine. Scholars who study networks ask questions like the following: Who are the most important people and collaborations in a network of overlapping committee memberships? What is a good movie-recommendation strategy in a social network? How did ideas spread over Twitter and other social media in the #Brexit debate, and how did this spread of ideas influence opinions and events? What role did the spread of misinformation, "fake news", and "alternative facts" play in the 2016 US presidential election, and how did this role change with evolving network structure as communication channels became more polarized? How can one use information from social interactions to improve strategies for predicting and preventing criminal activity? Which parts of a granular material are the least stable, and how should one measure this? How can one improve transportation systems and building layouts (e.g., the location of checkout tills and sale items in a supermarket) to ease traffic congestion? How can one control cascading failures in infrastructure or in financial networks? How should one measure the robustness of power grids to failures in transmission lines, and how should one design smart grids to ensure robustness? How do financial assets coevolve, and what are the best measurements to help predict (or at least characterize) their future evolution? How does the structure of an animal social network affect individual and collective behavior, and what types of interactions should one consider in such a network?

Given data in which connectivity plays a role, and assuming that one wishes to use tools from network science to help in one's analysis (that itself is not always obvious, and it is a modeling decision), one needs to decide what type of network description to use, analogously to deciding a level of description using other approaches (e.g., discrete particulate interactions versus a continuum (PDE) model). Just as one needs to use the correct conservation laws (and boundary conditions, initial conditions, and so on) in continuum models, it is crucial to choose an appropriate network representation for the problem at hand. If one studies the wrong network or asks questions that one's network representation cannot answer adequately, it is easy to end up with nonsense.

For example, one may possess a time-resolved data set representing social interactions. The most common approach in such a situation is to aggregate the data into a time-independent representation and study the resulting graphs and adjacency matrices with standard tools, but that can cause several problems. There are many choices for aggregation, the simplest of which is to simply count the number of interactions between each pair of entities and place those numbers in the weight matrix $\mathbf{W}$. If $w_{ij} = 2$, entities i and j interacted with each other twice during the monitored time window. Unfortunately, this type of aggregation ignores the "bursty" nature of dynamics in social systems and instead makes the (usually) incorrect assumption of Poissonian temporal interactions [28]. One can try to aggregate the temporal information in a more sophisticated way, but then one has to think very carefully about both the observations and sociological (or other) model of communication between individuals. Such aggregations of temporal data into time-independent representations also suffer problems related to concurrency and ordering of interactions, which is crucial for applications such as transmission of information and diseases (and thus for many problems of industrial interest), so one needs to go beyond the traditional tools associated with time-independent graphs. One can thus study "temporal networks" [31, 30] and either perform aggregations over multiple windows (which can either overlap or not) or perhaps not aggregate at all and consider the timeline of interactions to be the objects of interest — and indeed the study of temporal networks is a very active area of network science — but one still faces several problems (related to unresolved theoretical issues on which there is active research):

1. It is very far from clear how to generalize measures and approaches (e.g., "centrality" measures of node and edge importance, data clustering methods, and others) from time-independent networks to temporal networks.

2. These concepts can be generalized in many different ways, and it is an open issue as to which generalizations are better for which applications, problems, and data.

3. One has to consider the important issues of discrete versus continuous time and of interaction duration.

4. One must also consider the relative temporal scales of changes in network structure (weights and/or connections) and changes in the states of network nodes and edges (e.g., time-dependent traffic on city streets).

Thus, with temporal networks (and any other generalization of ordinary graphs) there is a modeling tradeoff: Should one collapse the data and use a simpler representation, possibly losing something vital or even obtaining a qualitatively incorrect answer in the process, to be able to use a better-developed and better-understood approach; or should one keep some of the salient information (surely one cannot keep all of it, because of limitations imposed by data size) and have to generalize the mathematical approach and perhaps make some missteps along the way? As with mathematical modeling (and very prominently indeed in

industry), the answer is that it depends on the problem and the question that one is asking. Ideally, one examines both approaches, because it is necessary develop a better understanding of which simplifications are acceptable.

Temporal dynamics is not the only type of complication in interaction data. For example, data can have "multilayer" structures, perhaps through the interaction of multiple subsystems or through the presence of multiple types of connections (e.g., there can be multiple communication channels or multiple modes of transportation) [34, 7], and one thus has to consider whether to use a monolayer or multilayer approach. As with temporal networks, one keeps more information with a multilayer approach, but it is challenging to generalize monolayer measures and methods, as different generalizations are appropriate in different situations. There are several other similar issues. Should one consider just network structure, a dynamical process on a network, or an "adaptive network" [57], in which the dynamics on top of networks are coupled to the dynamics of the network structure (e.g., a driver changes his/her route based on traffic conditions)? Should one include annotations (e.g., demographic data) on nodes and/or edges? Should one allow "hyperedges" to connect more than two types of nodes? Should one perhaps do all of these things (as well as others that we have not mentioned)? However, including everything yields an enormous mess that nobody knows how to study!

Further issues arise in the form and fidelity of data. Data may be inaccurate or missing (and "Big Data" is very far from the same thing as "good data" or "appropriate data"), and generalizing network structure to incorporate more features necessitates demanding reasonable measurements of more things. Thus, what data one can reliably collect (or obtain access to) will also influence the complexity of the chosen network representation. There is also the issue that most data do not come initially in the form of a network, or it may come in such a form but with difficulty in determining the weights (or even existence) of edges between entities.[2] In some situations (e.g., physical networks, such as in the study of traffic on road networks), there are often straightforward ways to measure edge weights. In a lot of data (e.g., from social or biological interactions), however, it is often much harder to reliably calculate the weight of a network edge from empirical data. For example, suppose that data arises in the form of pairwise similarities between entities. One can construct an adjacency matrix and thus a network, but is a network approach the best one (or even a good one) to use? Perhaps one should instead use data-reduction techniques from machine learning? To give an even more complicated scenario, one may possess coupled time series, so one can construct a (possibly time-dependent) set of similarities using one or more of many possible ways of measuring similarities between time series, and one thereby obtains a network (either a temporal one or time-independent one) to analyze. But one started with a set of coupled time series, so maybe one should use time-series approaches?

Another issue that arises in the study of networks involves the notion of "modeling" itself. An applied mathematical (or physical) model usually takes

[2] As Tom Petty and the Heartbreakers might sing, the weighting is the hardest part.

as a starting point a postulated mechanism of cause and effect. Statistical models, in contrast, take the data as the starting point and are therefore more (and sometimes purely) descriptive in nature. Statistical models are often very successful at indicating correlations (interpreted generally). Consequently, it is not surprising that the study of networks is more common among statisticians than among applied mathematicians. However, statistics is rightly cautious in deducing causation from correlation. Hence, a distinctive feature of "applied mathematical" network modeling is its linking of ideas and tools from statistics (which are necessary, given the coupled and high-dimensional nature of networked systems) with the desire for causal mechanisms. One possible desirable outcome is to derive some kind of equations of motion (perhaps high-dimensional ones), which one can try to simplify using some sort of mean-field theory, master equation, or other approximations [48, 53]. Unfortunately, this is often very difficult.

Another salient point, which we have seen Andrew Stuart point out [61] in the context of data assimilation (see [38] for an introduction to data assimilation) and which we borrow here for our discussion, is the level of verifiability (a kind of trust) of models in different domains and how that affects mathematical modeling and the interaction between data and models. At one extreme — i.e., most verifiable — lie physical models, in which one has a mechanistic model that, in many cases, one trusts fully because it is derived from fundamental physical principles that are supported by numerous repeatable experiments with very precise and accurate results. Somewhat less verifiable (or calibratable) are typical models from biology, which are often exploratory phenomenological models; nevertheless, there is still often a strong level of support and/or confidence in them, and they often incorporate causal mechanisms based on observation. Models from mathematical finance, as used by financial institutions, often also lie at this level; they can be well-calibrated to data, and they have a built-in causal mechanism (e.g., no arbitrage), but they also have some inbuilt ad hoc assumptions. Even trickier (in the context of verifiability) are many models in the social sciences, which are often based on sociological theory and/or thought experiments with much less direct observational support than the ones in biology or mathematical finance. Their role is often to illustrate or probe a postulated mechanism or process embedded in a larger and more complex context, but with little (or no) expectation of direct comparison with data. (See Section 3 for examples and discussion of mathematical modeling in the social sciences.) Some of these models also rely on assumptions, such as perfect rationality, whose validity is hotly contested. Finally, there are models, in fields such as commerce and many others, where one usually possesses only empirical data and simple agent-based or machine-learning models, but one aims to tackle large-scale and complex problems by simulation. An important question is the following: How does the interaction between modeling and data differ in these different scenarios? In a physical model, in which one believes (at least approximately) the equations of motion, data may simply yield initial conditions, boundary conditions, and parameter values. At the other end of the spectrum, for many complex systems, one does not have equations, and a model

may be entirely data-driven or even purely descriptive. See [37] for additional discussion of data-driven modeling in complex systems (for scenarios with various levels of trust). It is important to emphasize that the location of a particular model on this spectrum depends strongly on the availability of widely-accepted basic principles, but it is also fair to comment that each discipline has its own prejudices in favor of some kinds of model and against others.

Many problems (especially ones that involve complex systems) come in a form in which network science can help. A key challenge in network science — and a major difference with other areas of industrial and applied mathematics — is that the off-the-shelf tools are much less developed in network science than they are in other areas. It is far from clear how to generalize tools and methods from graphs to more complicated types of networks (there remains a wealth of open problems even when studying graphs), and such efforts are the most active part of network science. When faced with a problem from industry (or elsewhere), there is a tension between simplifying the problem and using off-the-shelf tools and trying to develop the new mathematical and computational tools that are necessary to examine the problem in a more detailed and perhaps more appropriate way.

3 Applying Mathematics to the Social Sciences

An area that helps set the stage for the importance of networks in industrial applications is the application of mathematics, and networks especially, in the social sciences. The use of mathematical approaches to social phenomena is very old (and much older than appreciated) [63, 15]. More recently, the wealth of social data (e.g., from social media such as Twitter and Faceook, and directly from companies in forms such as mobile-phone data, shopping data, human movement data, and others) has brought social science to the center of the "Big Data" explosion. In contrast to much smaller data, collected traditionally in forms such as surveys, the data deluge has led to the formation of subjects such as "computational social science" [39, 56]. This has placed social science in a transition period in which an increasing number of researchers who are trained in subjects such as computer science, physics, and mathematics are trying to apply their techniques to social systems [65], which are undergoing a revolution in our ability to predict and explain their dynamics [29].

3.1 Example: Influence and Opinions on Networks

To give a concrete example of mathematical modeling in the social sciences, we discuss ideas related to opinions and social influence on networks [53, 8]. One of the best-known, albeit in many ways rather naive, model for social influence on networks is the Watts Threshold Model (WTM) [64], which models the adoption of a product, idea, innovation, or meme in a social network. Independently, the WTM was studied slightly earlier by Valente [62], who incorporated network structure into an influence update rule that Granovetter proposed in the

1970s [23]. The WTM is also a generalization of a process known as "bootstrap percolation [10]. In the WTM model, and other models of social influence, it is interesting to ask if an idea takes root (e.g., if a product "goals viral" and reaches a large fraction of individuals in a network), how long it takes to take root, where in the network it is optimal to seed a new product, and so on. There are also numerous other types of opinion models, often in the form of "voter models", in which people ask whether consensus occurs, how long it takes to occur, how many different groups of different opinions are present at equilibrium, and so on [53, 8]. These results depend both on network structure and on the rules that govern how node (or edge) states are updated, and different qualitative results can result from rather specific, seemingly-minor choices (e.g., drawing an edge uniformly at random in a voter model versus drawing a node uniformly at random in a voter model and then picking one of its incident edges) [53, 41, 58]. The fact that such choices can have drastic effects on qualitative dynamics is a major modeling issue, especially if one starts with an ill-defined, woolly problem, as is almost invariably the case when collaborating with industry.

Suppose that one has a network, which can arise either from empirical data or from a synthetic construction (perhaps as an instantiation of a random-graph model). For simplicity, let's suppose that this network is represented by a graph and is both unweighted and undirected. The WTM is a binary-state "threshold model" [21]. In such models, each node i has a threshold R_i (which, in most models, is time-independent) that is drawn from some distribution. At any given time, each node can be in one of two states: 0 (inactive, not adopted, not infected, etc.) or 1 (active, adopted, infected, etc.). The states of the nodes change in time according to an update rule, and one can update nodes either synchronously or asynchronously. (The latter leads to approximations in terms of continuous-time dynamical systems.) When updating the state of a node in the WTM, one compares the node's fraction m_i/k_i of infected neighbors (where m_i is the number of infected neighbors and k_i is the degree of node i) to the node's threshold R_i. If node i is inactive, it then becomes active (i.e., it switches to state 1) if $m_i/k_i \geq R_i$; otherwise, its state remains unchanged. One should think of the quantity m_i/k_i as a peer pressure, and one way to generalize the WTM is by calculating peer pressure in different ways. In a model introduced in [44], for example, there are three states: nodes can be passive, active, or "hyper-active", where the last category of nodes, which could represent leaders in a mass movement, exert more influence than nodes who are merely active.

Threshold models are rather simplistic, and it is natural to ask whether there exist real-life scenarios in which such models are appropriate for explaining empirical observations [53, 59]. Although a binary decision process on a network is a gross oversimplification of reality, it can already capture two very important features [49]: interdependence (an entity's behavior depends on the behavior of other entities) and heterogeneity (differences in behavior are reflected in the distribution of thresholds). Typically, some seed fraction $\rho(0)$ of nodes is assigned to the active state, although that is not always true (e.g., when $R_i < 0$ for some nodes i). Depending on the problem under study, one can choose the

initially active nodes via some random process (typically, uniformly at random) or with complete malice and forethought. For the latter, for example, one can imagine planting a rumor with specific nodes in a network, or perhaps this is a seed node that is trying to spread misinformation, "fake news", or "alternative facts".

The study of social influence on networks is important for both commercial and societal considerations. What ideas (or products or memes) will become viral, how should one model the spread of such ideas (recall our comment above that very small differences in the model can result in rather different qualitative results), and how should one measure virality in the first place? These are important open problems, and they are relevant for numerous governmental and industrial stakeholders. In measuring virality, for example, an adoption "tree" (technically, it is a directed acyclic graph [48]) can have many different shapes, and one needs to measure both lengths (e.g., the length of the longest adoption chain) and widths (e.g., the largest number of children of any given node in the tree), as, for a fixed number of adopters, long-and-thin trees can predict very different qualitative features (in particular, the eventual number of adopters) for subsequent adoption than short-and-fat trees [22]. The time distributions and recurrence patterns of adoptions are also very important [12].

There are many commercial and governmental applications of social influence and opinion dynamics on networks. An effective model of a marketing campaign — whether to promote a product (or an idea or desired behavior, such as a healthy diet), to perform targeted and personalized advertising, or to counteract the spread of misinformation — requires one to understand how network structure can affect network dynamics and how different dynamical processes (even ones that cosmetically seem rather similar to each other) can exhibit qualitatively different behaviors. These are active research areas with numerous fascinating theoretical issues (in mathematics, social science, human behavior, economics, and more), methodological issues (for both analytical theories and computation), and commercial, governmental, and societal applications. In the commercial sector, for example, research on human opinion, behavior, and influence can play a significant role in personalized coupons. If, as sociological theory suggests, "you are who your friends are" (as social network structure arises in large part from homophily) [63], one can imagine that coupon-program design should include information about what one's neighbors in a social network are buying.

3.2 Ethics

Privacy is one of several crucial ethical issues raised by the analysis of personal data [25]. One can use network data to de-anomyize people, especially when there is data about many of the same people in different social networks [3, 46], and some personal data (e.g., medical records) can influence insurance premiums or other important items.

Most mathematical scientists do not have appropriate ethical training for the era of Big Data, and it is necessary that such training be built into their

education. Much of the social data used in collaborations with industry make these issues arise rather prominently, so this will become an increasingly important aspect of industrial mathematics with social data — and especially with network data, as the ties between people can play a major role in invasion of privacy and removal of anonymity. Several years ago, we (MAP, with help from SDH) set up a procedure for studying human data for the Mathematical Institute at University of Oxford (the current version is available at [43]), and we would like to see this kind of approach becoming standard.

While privacy may be the most obvious ethical dilemma facing researchers who analyze personal data, other ethical issues are also present, and applied mathematicians may encounter them in their research. To highlight recent salient examples for which ethical considerations are important, one team of scholars use of Facebook data for an experiment about emotional contagions wrought controversy concerning the manipulation of user emotions based on experiments with adjusting user Facebook feeds [36, 11], and the UK company Cambridge Analytica used data-driven targeted advertising (via psychological profiles from social media activity) for the Trump US presidential campaign [24].

4 Connecting with Industry

During the last decade, we have cosupervised several doctoral students at University of Oxford jointly with industrial partners, initially through EPSRC Collaborative (Industrial CASE) awards and more recently through programs such as the Mathematical Institute's Centre for Doctoral Training (CDT) in Industrially Focused Mathematical Modelling (InFoMM) [42]. All together, this covers half a dozen students: four in collaboration with the investment bank HSBC, one with the supermarket company Tesco, and one with the customer-science company dunnhumby.

Our collaboration on network science started with our students who worked on problems in collaboration with HSBC, and (*very importantly*) it has led directly both to results of interest to stakeholders (HSBC and their clients) and to us and the applied mathematics and network-science communities more broadly. As these projects illustrate, the applications and the mathematics are intertwined, as each drives the other in a crucial way.

Early work on time-dependent correlations of financial assets [18, 16, 17] included ideas from network science and random-matrix theory [18] and led to HSBC's "risk-on, risk-off" (RORO) strategy. RORO has been featured prominently in financial circles (including as a tag in the *Financial Times* blog); see, e.g., [32].

This work with HSBC helped pave the way for our more recent work on financial assets, multilayer networks, and current applications to consumer–product purchasing networks, and in parallel it stimulated theoretical analysis of time-dependent networks. The starting point was the study of "community detection" [54, 19], an approach to network clustering in which, in some hope-

fully optimal way, one seeks to algorithmically find dense sets of nodes that are connected sparsely to other dense sets of nodes. In the HSBC-related work in [16, 17], led by our doctoral student Dan Fenn, communities were detected by optimizing the "modularity" objective function [48, 47]. A measure of similarities in the time series of the exchange rates, based on a time-window aggregation, was used to compare observed connections in a network with "random ones" in a null network constructed from a random-graph model. This was done separately in each window, but the windows were connected sequentially, because there is a temporal overlap between consecutive windows.

Informed by this work (and also work in other applications, such as legislative cosponsorship and voting networks [67, 66]), MAP and collaborators developed a more principled approach to community detection in time-dependent networks with discrete temporal "layers" that can represent interactions at one time point or over some period of activity [45]. In this approach, one incorporates the "contiguity" of the layers using "interlayer" edges between nodes in different layers. Mucha et al. [45] derived a generalization of the modularity objective function for the resulting "multilayer" network [34], and they maximized it to assign node-layer pairs to communities. Thus, an entity can be assigned to different communities in different time periods, and one can study the evolution of community structure over time. Multilayer modularity can also be used for other situations, such as "multiplex" networks, in which there exist multiple different types of edges.

The work of [45] left open numerous questions about multilayer networks in general and about community detection in such networks in particular. Our student Marya Bazzi, also funded by a CASE award in collaboration with HSBC, revisited some of the applications and data sets studied by Dan Fenn years earlier but now used the more sophisticated multilayer-network techniques, which had been developed and advanced during that time, to examine those phenomena in a more sophisticated manner [6]. (Following the 2010 publication of [45], the study of multilayer networks has become one of the most prominent areas of network science [34, 7]; and several papers, such as [50, 14] and others, have proposed different approaches for studying communities and other mesoscale structures in such networks.) In addition to affirming results from the work of Fenn et al. [18, 16, 17] (such as structural changes in the networks following the Lehman Brothers bankruptcy in 2008), she also found subtler structural changes that merit further exploration. As the focus of the work shifted from the application to theory, Marya made several theoretical and methodological advances, first in [6] and then (jointly with us, fellow University of Oxford doctoral student Lucas Jeub, and our collaborator Alex Arenas) in [5]. Advancements in [6] include the careful distinction between null models and null networks in modularity maximization, proofs of crucial conceptual ideas in multilayer networks (e.g., that the limit of zero interlayer coupling is a singular one), and toy examples that set the stage for our recent systematic development of flexible benchmark models for multilayer networks in [5]. These benchmarks are generative models that allow wide variety of correlations across different layers — a feature that is very important for real multilayer networks.

Our current work in collaboration with supermarket and customer-science companies partly builds on the above insights on clustering in networks and partly moves entirely in new directions. In one project, we are clustering shopping data and trying to incorporate constraints and metadata that are appropriate for that application. In another, we are developing a generative model for the shopping trajectories of people in supermarkets.

The project of our doctoral student Roxana Pamfil lies within this strand of research but also considers a new application, which has structural constraints that differ in important ways. Our work with Roxana is in collaboration with dunnhumby, a customer-science company and subsidiary of the supermarket chain Tesco. Roxana is analyzing data from anonymized consumers and the products that they purchase in their shopping baskets (from many Tesco stores in the United Kingdom). From these data, we construct bipartite (i.e., two-mode) networks in which consumers are adjacent to purchased products. See Fig. 1 for a schematic. The bipartite structure needs to be incorporated into methods for clustering the data. Determining the edge weights also requires considerable care. For example, one can use so-called "item penetration" (the fraction of all of the items bought by customer c that were product p), "basket penetration" (the fraction of all baskets of customer c that include product p), or something else. Different choices can yield qualitatively different results, and a key challenge is to determine precisely which weights are most appropriate for which questions and which of them give the most robust results. Another crucial consideration is the time windows over which to examine purchases and the fact that there is data from multiple stores (including different store formats, with stores located in different regions of the UK).

Because we also have access to time-resolved data, which are collected from several different stores (including multiple different store formats), and include various shopper metadata, we hope eventually to study product–purchase networks that are both temporal and multiplex (i.e., that have two different multilayer features), an important challenge in the study of multilayer networks. In practice, this last challenge will likely be beyond the scope of Roxana's thesis and we will use only a subset of these data in our analysis, and we still need to decide exactly what the subset will be. We also have access to product descriptions at different hierarchical "levels" (e.g., organic milk versus a particular type of organic milk), opening the door for multilevel modeling with interlayer edges that represent inclusion relationships and induced intralayer edges whose existence can be inferred on a different layer of a multilayer network. (For example, if a consumer bought a particular type of organic milk, then he/she necessarily purchased organic milk more generally.) For examples of networks with various complications, see Fig. 2.

An important aspect of our project with dunnhumby is clustering with mesoscale structures aside from the "assortative" ones (which correspond to adjacency matrices that look dense in the main block diagonal) that are typified by community structure [19]. One way to do this in a statistically principled way is to use "stochastic block models", in which one specifies a block structure and tries to find the clustering that best fits that structure [19]. With

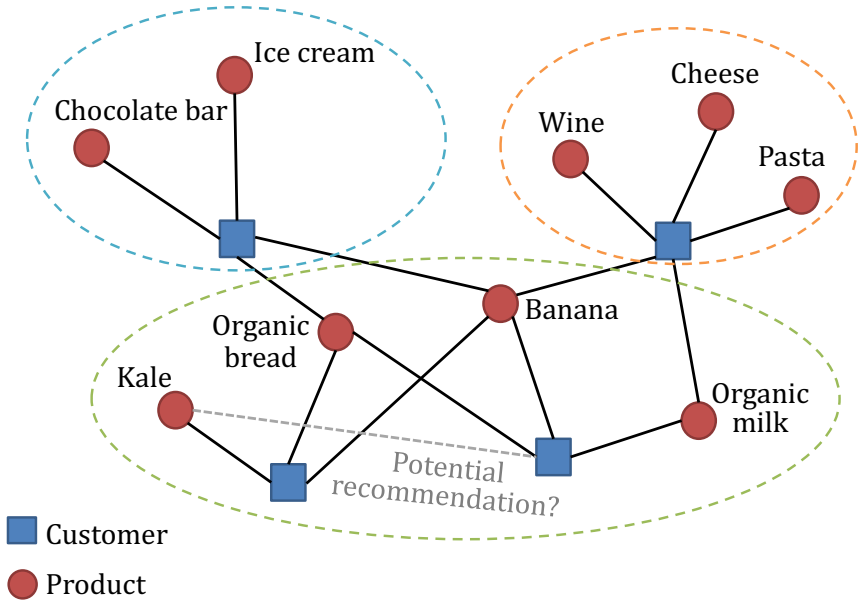

Figure 1: A schematic bipartite (i.e., two-mode) network in which consumers are connected to the products that they purchase. The blue squares are consumers, and the red circles are the products that they have purchased. The dashed ellipses enclose nodes that have been assigned to the same community using a clustering technique. The gray dashed edge in the lower part of the figure is a potential recommendation that one might give to a consumer, as there is a node that is not yet purchasing kale but who has been assigned to the same community as kale. [This figure is a slight modification of a figure that was created by Roxana Pamfil.]

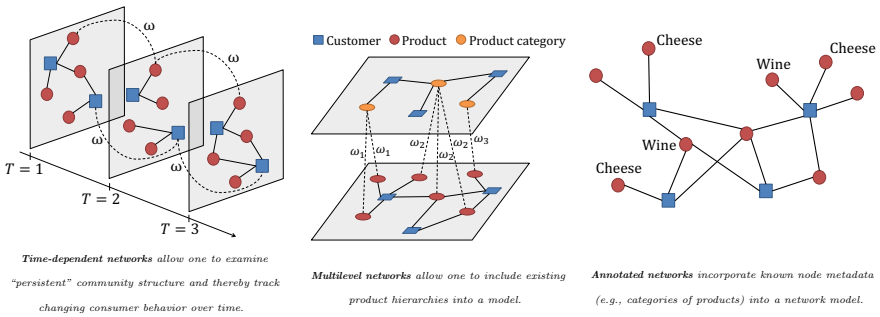

Figure 2: Schematic of network structures that we may study as part of our work on product–purchase networks. (Left) Time-dependent network, encoded with a multilayer representation, with time layers that indicate purchases by consumers (blue squares) of products (red circles) in supermarkets. The interlayer edge weights ω encode dependencies across time layers. Determining values for ω, ideally from data or as an output of analysis, is an open problem in the study of multilayer networks. (Center) A multilevel network, in which different layers represent different hierarchical levels in product descriptions. Note that the interlayer edge weights ω_i can be heterogeneous, as is also the case for multilayer networks more generally. (Right) Annotated networks, in which we use product data as product-node labels. [This figure is a slight modification of a figure that was created originally by Roxana Pamfil.]

this kind of perspective, statistical inference takes center stage, and statistical model selection — e.g., between different block structures or between whether or not one allows overlapping communities — becomes a key consideration [51]. In a multilayer setting, as we showed in [5], one can incorporate interlayer dependencies directly into generative network models (such as stochastic block models) in a convenient way, and we are actively using these ideas for clustering in consumer–product purchasing networks.

Roxana's still-evolving project is a good illustration of the way in which model choices, analysis, and simulation interact with data in an iterative and question-driven way. A crucial point to stress once again is the benefit to both industrial and academic partners. In this project, the information in the data and the structural constraints of the application yield multilayer networks with different structures and metadata (and different sparsity patterns in the network connections) than what has been analyzed previously. To further develop the theory of multilayer networks, which is one of our primary scientific interests and is one of the most active areas in network science, it is necessary to consider diverse structures, applications, and ensuing challenges. Otherwise, one risks developing a biased theory that hasn't been tested adequately on those types of structures. For our industrial partners, Roxana's project will result in a thorough understanding of different types of edge weightings between consumers and the products that they purchase, how to categorize different types of customers, and the development of strategies for product recommendations and personalized coupons (through reward cards, which are also helpful for gathering data). It will also account for geographic variabilities in how people shop and for large-scale changes in customer preferences over time. Eventually, we would like to combine these insights with social-media data (e.g., with recommendations that are also influenced by friends' purchases), though that will of course involve very serious ethical considerations regarding what research in that direction is appropriate.

Our latest doctoral student, Fabian Ying, who is working on a project in collaboration with Tesco. His will thus also examine some of their data, though his project — examining human mobility inside supermarkets — departs rather markedly from the ones that we described above. The study of mobility in supermarkets is important for several questions (see [55] for a recent study), including the following: How do customers shop and navigate within a supermarket? What is the best store layout to reduce congestion? Where in the stores should the promotional items be placed? For our industrial partner, these questions are of course related to one of their major questions: How to maximize revenue?

Thanks to recent technology, customer trolleys (which are known as "shopping carts" in some locales) can now be equipped with RFID chips, allowing large-scale, anonymized, and non-invasive mapping of customer journeys. This makes it possible, in principle (though we are currently trying to overcome challenges from the messy nature of the data), to conduct rigorous investigations of key questions like the ones above. As part of a trial for this trolley-tracking technology, Tesco has collected trolley location data in a couple of of its stores

and has provided us with several days of data.

Our approach in this project is to build a simple mathematical model for customer movements within a store. We are using insights from existing models for human mobility, which have been a major research topic during the last decade thanks to the recent availability and abundance of human and animal mobility data (see, e.g., [9] and references therein). However, our project entails analysis of rather different temporal and spatial scales from existing studies. This, in turn, necessitates the development of new mathematical models, which include inherently interesting mathematics, will lead to insights on human mobility (e.g., shopping trajectories), and will help Tesco develop a strategy to answer their key questions (e.g., reducing congestion).

Based on the spatial resolution of the RFID chips, we model each store as a network whose nodes are different regions of the stores, with edges present whenever two regions are next to each other. There are also specific entrance and exit points (the tills), and these networks are directed. Our model, which we are still developing, is based on random walks on networks [41], although one cannot simply use the most naive type of random walk, as it is crucial to incorporate additional features, which can include shopping lists, impulse purchases, a probability to stay in the same place and continue shopping, the effects of local congestion, and so on. Our current model includes a shopping list (and one can imagine shoppers with heterogeneous lists), a probability of taking a detour, and an attractiveness value for each zone that influences random-walk probabilities. The spatial embeddedness of a supermarket network is also an important consideration, as the embeddedness of spatial networks induces special structural features that make them different from other networks [4]. Right now, we are studying a random-walk model on small, simple networks; and the fact that the real networks in this case are embedded in two dimensions influences our choice of which toy networks to consider. An exciting aspect of Fabian's project is that we are revisiting classical results from stochastic modeling (e.g., queuing theory), dusting them off, and adapting them for the network age.

Fabian's research is in its early stages. We are developing a toy model in close consultation with Tesco, and we are comparing the human trajectories that we see using our nascent (and still changing) models on both synthetic networks and empirical networks from the RFID data. The data is extremely noisy, so at the moment this is a major challenge. Potential future avenues include incorporating insights from recommender systems (e.g., using mobility and behavioral economics) to influence movement and avoid congestion. We also hope to compare the results from our toy models to numerical computations using more realistic, agent-based models.

5 Conclusions

Network science is playing a large — and increasing — role in industrial problems. Many problems, and associated data, have a natural network structure;

and the study of networks and other discrete structures is rapidly becoming a core area of applied mathematics alongside traditional continuum approaches [52]. (The study of networks has, of course, been prominent in subjects like graph theory for many decades.) The traditional and very successful "physical applied mathematics" philosophy is just as relevant in network modeling as it is in more traditional applied-mathematics topics, but there are also many fascinating, important, and often rather difficult challenges: (1) the field of network science is much less mature than topics such as partial differential equations and asymptotic analysis, and this leads both to the development of new methodologies from industrial (and other application-oriented) problems and to less clarity in how and what level of description to use to attack those problems; (2) because networks are high-dimensional and the interactions between many entities play a prominent role in network analysis, network modelers should become comfortable not only with traditional mechanistic modeling, a longstanding expertise of applied mathematicians, but also with ideas such as statistical modeling (which is more descriptive) and uncertainty quantification; (3) the large scale of networked systems poses challenges for scientific computation, especially given not only large static data sets but also real-time computations with data streams; (4) missing and incomplete data (and data cleaning) provide significant guidance (and limitations!) that affect not only what calculations are reasonable but also the level of detail that one may wish to use in a network description; and (5) the vast and increasing use of human data poses significant ethical issues (e.g., data privacy) that departs rather markedly from, say, the data that arises from fluid mixing in chocolate and other traditional industrial applications that motivate mathematical studies.

The study of networks is a core part — and, we would argue, one of the most important parts — of the mathematics for the modern economy, as it relates very strongly both to the types of problems and to the types of data that arise in it. Network modeling also has deep connections to data analysis, data science, and Big Data, but it is more than that: network science incorporates modeling tenets from both physical applied mathematics and statistics, and it seeks to marry them together. Mathematics departments need to develop strength in network modeling, and those efforts need to include the study of problems with close ties to problems that arise in the industrial, commercial, and governmental sectors. This will help solve important societal problems and simultaneously lead to the develop of new mathematical and computational techniques and insights.

Acknowledgements, Data Ethics, and Other Details

Ethics: This paper includes a discussion of important ethical issues, especially with respect to data privacy, but it does not use any human data and thus this paper itself satisfies all ethical requirements.

Data access: Ironically, this paper contains no data.

Author contributions: Both authors contributed substantially; MAP wrote the

first draft and is giving the associated talk at the Royal Society.

Competing interests: We have no competing interests.

Funding: Our doctoral students Fabian Ying and Roxana Pamfil were funded by the EPSRC Centre For Doctoral Training in Industrially Focused Mathematical Modelling (EP/L015803/1) in collaboration with Tesco and dunnhumby respectively. Our doctoral students Daniel Fenn and Marya Bazzi were funded by the EPSRC through CASE studentship awards.

Acknowledgements: We thank John Ockendon for the invitation to contribute this article. We thank Robert Armstrong, Mariano Beguerisse Díaz, Jeremy Bradley, Peter Grindrod, Roxana Pamfil, Rosie Prior, and Fabian Ying for helpful comments. We also thank Roxana Pamfil for producing figures for this paper. Finally, we thank Andrew Stuart for his inspiring slide and discussion, which we adapted for one of our paragraphs in this paper.

Disclaimer: We have no disclaimers.

References

[1] Mathematics in Industry: Information Service. Available at `http://www.maths-in-industry.org`, 2017.

[2] P. J. Aston, A. J. Mulholland, and K. M. M. Tant, editors. UK Success Stories in Industrial Mathematics. Springer-Verlag, 2016.

[3] L. Backstrom, C. Dwork, and J. Kleinberg. Wherefore art thou r3579x? Anonymized social networks, hidden patterns, and structural steganography. In Proceedings of the 16th International Conference on World Wide Web (WWW '07), pages 181–190, Banff, Alberta, Canada, 2007. ACM.

[4] M. Barthélemy. Spatial networks. Physics Reports, 499(1):1–101, 2011.

[5] M. Bazzi, L. G. S. Jeub, A. Arenas, S. D. Howison, and M. A. Porter. Generative benchmark models for mesoscale structures in multilayer networks. arXiv:1608.06196, 2016.

[6] M. Bazzi, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison. Community detection in temporal multilayer networks, with an application to correlation networks. Multiscale Modeling and Simulation: A SIAM Interdisciplinary Journal, 14(1):1–41, 2016.

[7] S. Boccaletti, G. Bianconi, R. Criado, C.I. Del Genio, J. Gómez-Gardenes, M. Romance, I. Sendina-Nadal, Z. Wang, and M. Zanin. The structure and dynamics of multilayer networks. Physics Reports, 544(1):1–122, 2014.

[8] C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. Reviews of Modern Physics, 81:591–646, 2009.

[9] S. Çolak, A. Lima, and M. C. González. Understanding congested travel in urban areas. Nature Communications, 7:10793, 2016.

[10] J. Chalupa, P. L. Leath, and G. R. Reich. Bootstrap percolation on a Bethe lattice. Journal of Physics C: Solid State Physics, 12(1):L31–L35, 1979.

[11] C. Chambers. Facebook fiasco: Was Cornell's study of 'emotional contagion' an ethics breach? Available at `https://www.theguardian.com/science/head-quarters/2014/jul/01/facebook-cornell-study-emotional-contagion-ethics-breach`, 2014.

[12] J. Cheng, L. A. Adamic, J. M. Kleinberg, and J. Leskovec. Do cascades recur? In Proceedings of the 25th International Conference on World Wide Web, WWW '16, pages 671–681, Republic and Canton of Geneva, Switzerland, 2016. International World Wide Web Conferences Steering Committee.

[13] W. J. Cook. In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation. Princeton University Press, Princeton, USA, 2014.

[14] M. De Domenico, A. Lancichinetti, A. Arenas, and M. Rosvall. Identifying modular flows on multilayer networks reveals highly overlapping organization in interconnected systems. Physical Review X, 5:011027, 2015.

[15] I. de Sola Pool and M. Kochen. Contacts and influence. Social Networks, 1(1), 1978.

[16] D. J. Fenn, M. A. Porter, M. McDonald, S. Williams, N. F. Johnson, and N. S. Jones. Dynamic communities in multichannel data: An application to the foreign exchange market during the 2007–2008 credit crisis. Chaos, 19(3):033119, 2009.

[17] D. J. Fenn, M. A. Porter, P. J. Mucha, M. McDonald, S. Williams, N. F. Johnson, and N. S. Jones. Dynamical clustering of exchange rates. Quantitative Finance, 12(10):1493–1520, 2012.

[18] D. J. Fenn, M. A. Porter, S. Williams, M. McDonald, N. F. Johnson, and N. S. Jones. Temporal evolution of financial market correlations. Physical Review E, 84(2):026109, 2011.

[19] S. Fortunato and D. Hric. Community detection in networks: A user guide. Physics Reports, 659:1–44, 2016.

[20] A. C. Fowler. Mathematical Models in the Applied Sciences. Cambridge University Press, Cambridge, UK, 1997.

[21] J. P. Gleeson. Binary-state dynamics on complex networks: Pair approximation and beyond. Physical Review X, 3:021004, 2013.

[22] S. Goel, A. Anderson, J. Hofman, and D. J. Watts. The structural virality of online diffusion. Management Science, 62(1):180–196, 2016.

[23] M. Granovetter. Threshold models of collective behavior. The American Journal of Sociology, 83(6):1420–1443, 1978.

[24] H. Grassegger and M. Krogerus. The data that turned the world upside down. Available at `http://motherboard.vice.com/read/big-data-cambridge-analytica-brexit-trump`, 2017.

[25] P. Grindrod. Beyond privacy and exposure: ethical issues within citizen-facing analytics. Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, 374(2083), 2016.

[26] P. Grindrod, D. J. Higham, and P. Laflin. The graph whisperers. In P. J. Aston, A. J. Mulholland, and K. M. M. Tant, editors, UK Success Stories in Industrial Mathematics, pages 271–279. Springer, 2016.

[27] D. J. Higham, M. Batty, L. M. A. Bettencourt, D. V. Greetham, and P. Grindrod. An overview of city analytics. Royal Society Open Science, 4(2), 2017.

[28] T. Hoffmann, M. A. Porter, and R. Lambiotte. Generalized master equations for non-Poisson dynamics on networks. Physical Review E, 86:046102, 2012.

[29] J. M. Hofman, A. Sharma, and D. J. Watts. Prediction and explanation in social systems. Science, 355(6324):486–488, 2017.

[30] P. Holme. Modern temporal network theory: A colloquium. European Physical Journal B, 88(9):234, 2015.

[31] P. Holme and J. Saramäki. Temporal networks. Physics Reports, 519(3):97–125, 2012.

[32] HSBC FX Quantitative Strategy Group. HSBC: Risk on–risk off: The full story. Currency Quant Special, HSBC Global Research, November 2010.

[33] M. O. Jackson. Social and Economic Networks. Princeton University Press, 2010.

[34] M Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter. Multilayer networks. Journal of Complex Networks, 2(3):203–271, August 2014.

[35] E. D. Kolaczyk. Statistical Analysis of Network Data: Methods and Notes. Springer-Verlag, 2009.

[36] A. D. I. Kramer, J. E. Guillory, and J. T. Hancock. Experimental evidence of massive-scale emotional contagion through social networks. Proceedings of the National Academy of Sciences of the United States of America, 111(24):8788–8790, 2014.

[37] J. N. Kutz. Data-Driven Modeling & Scientific Computation: Methods for Complex Systems & Big Data. Oxford University Press, 2013.

[38] K. J. H. Law, A. M. Stuart, and K.C. Zygalakis. Data Assimilation: A Mathematical Introduction. Springer-Verlag, 2015.

[39] D. Lazer, A. Pentland, L. Adamic, S. Aral, A.-L. Barabási, D. Brewer, N. Christakis, N. Contractor, J. H. Fowler, M. Gutmann, T. Jebara, G. King, M. Macy, D. Roy, and M. Van Alstyne. Computational social science. Science, 323(5915):721–723, 2009.

[40] C. C. Lin and L. A. Segal. Mathematics Applied to Deterministic Problems in the Natural Sciences. Classics in Applied Mathematics (Book 1). Society for Industrial and Applied Mathematics, Philadelphia, USA, 1988.

[41] N. Masuda, M. A. Porter, and R. Lambiotte. Random walks and diffusion on networks. arXiv:1612.03281, 2016.

[42] University of Oxford Mathematical Institute. Industrially focused mathematical modelling (epsrc centre for doctoral training). Available at `https://www.maths.ox.ac.uk/study-here/postgraduate-study/industrially-focused-mathematical-modelling-epsrc-cdt` (accessed 29 January 2017).

[43] University of Oxford Mathematical Institute. Research using data involving humans. Available at `https://www.maths.ox.ac.uk/members/policies/data-protection/research-using-data-involving-humans` (accessed 29 January 2017).

[44] S. Melnik, J. A. Ward, J. P. Gleeson, and M. A. Porter. Multi-stage complex contagions. Chaos, 23(1):013124, 2013.

[45] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela. Community structure in time-dependent, multiscale, and multiplex networks. Science, 328(5980):876–878, 2010.

[46] A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets. In Proceedings of the 2008 IEEE Symposium on Security and Privacy (SP '08), pages 111–125, Washington, DC, USA, 2008. IEEE Computer Society.

[47] M. E. J. Newman. Finding community structure in networks using the eigenvectors of matrices. Physical Review E, 74:036104, 2006.

[48] M. E. J. Newman. Networks: An Introduction. Oxford University Press, 2010.

[49] P. Oliver, G. Marwell, and R. Teixeira. A theory of the critical mass. I. Interdependence, group heterogeneity, and the production of collective action. American Journal of Sociology, 91(3):522–556, 1985.

[50] T. P. Peixoto. Inferring the mesoscale structure of layered, edge-valued, and time-varying networks. Physical Review E, 92(4):042807, 2015.

[51] T. P. Peixoto. Model selection and hypothesis testing for large-scale network models with overlapping groups. Physical Review X, 5(1):011033, 2015.

[52] M. A. Porter and G. Bianconi. Network analysis and modelling: Special issue of *European Journal of Applied Mathematics*. European Journal of Applied Mathematics, 27(6):807–811, 2016.

[53] M. A. Porter and J. P. Gleeson. Dynamical Systems on Networks: A Tutorial. Springer-Verlag (Vol. 4 in "Frontiers in Applied Dynamical Systems: Reviews and Tutorials"), Heidelberg, Germany, 2016.

[54] M. A. Porter, J.-P. Onnela, and P. J. Mucha. Communities in networks. Notices of the American Mathematical Society, 56:1082–1097, 1164–1166, 2009.

[55] P. S. Riefer, R. Prior, N. Blair, G. Pavey, and B. C. Love. Coherency-maximizing exploration in the supermarket. Nature Human Behaviour, 1:0017, 2017.

[56] M. Salganik. Computational Social Science: Social Research in the Digital Age. Available at `http://www.princeton.edu/~mjs3/soc596_f2016/` (accessed 29 January 2017), Fall 2016.

[57] H. Sayama, I. Pestov, J. Schmidt, B. J. Bush, C. Wong, J. Yamanoi, and T. Gross. Modeling complex systems with adaptive networks. Computers & Mathematics with Applications, 65(10):1645–1664, 2013.

[58] V. Sood, T. Antal, and S. Redner. Voter models on heterogeneous networks. Physical Review E, 77(4):041121, 2008.

[59] B. State and L. Adamic. The diffusion of support in an online social movement: Evidence from the adoption of equal-sign profile pictures. In Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW '15, pages 1741–1750, New York, NY, USA, 2015. ACM.

[60] S. H. Strogatz. Exploring complex networks. Nature, 410:268–276, March 8 2001.

[61] A. M. Stuart. Private Communication, 2017.

[62] T. W. Valente. Network Models of the Diffusion of Innovations. Hampton Press, 1995.

[63] S. Wasserman and K. Faust. Social Network Analysis: Methods and Applications. Cambridge University Press, Cambridge, UK, 1994.

[64] D. J. Watts. A simple model of global cascades on random networks. Proceedings of the National Academy of Sciences of the United States of America, 99(9):5766–5771, 2002.

[65] D. J. Watts. Should social science be more solution-oriented? Nature Human Behaviour, 1:0015, 2017.

[66] A. S. Waugh, L. Pei, J. H. Fowler, P. J. Mucha, and M. A. Porter. Party polarization in Congress: A network science approach. arXiv:0907.3509, 2012.

[67] Y. Zhang, A. J. Friend, A. L. Traud, M. A. Porter, J. H. Fowler, and P. J. Mucha. Community structure in Congressional cosponsorship networks. Physica A, 387:1705–1712, 2008.